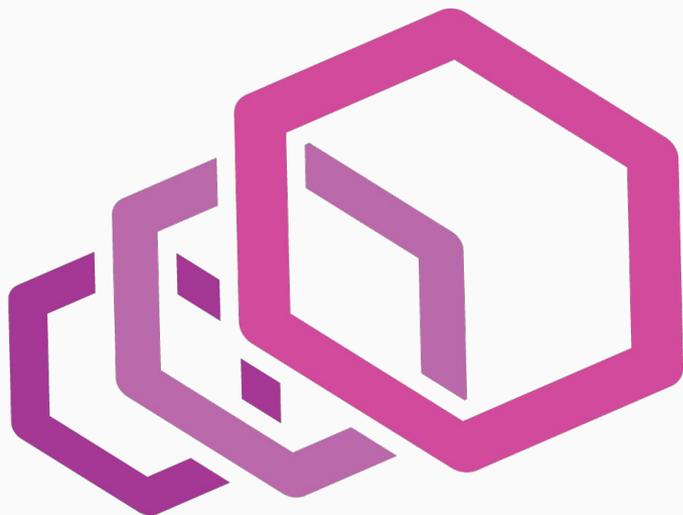


envoy

## The mechanics of deploying Envoy at Lyft

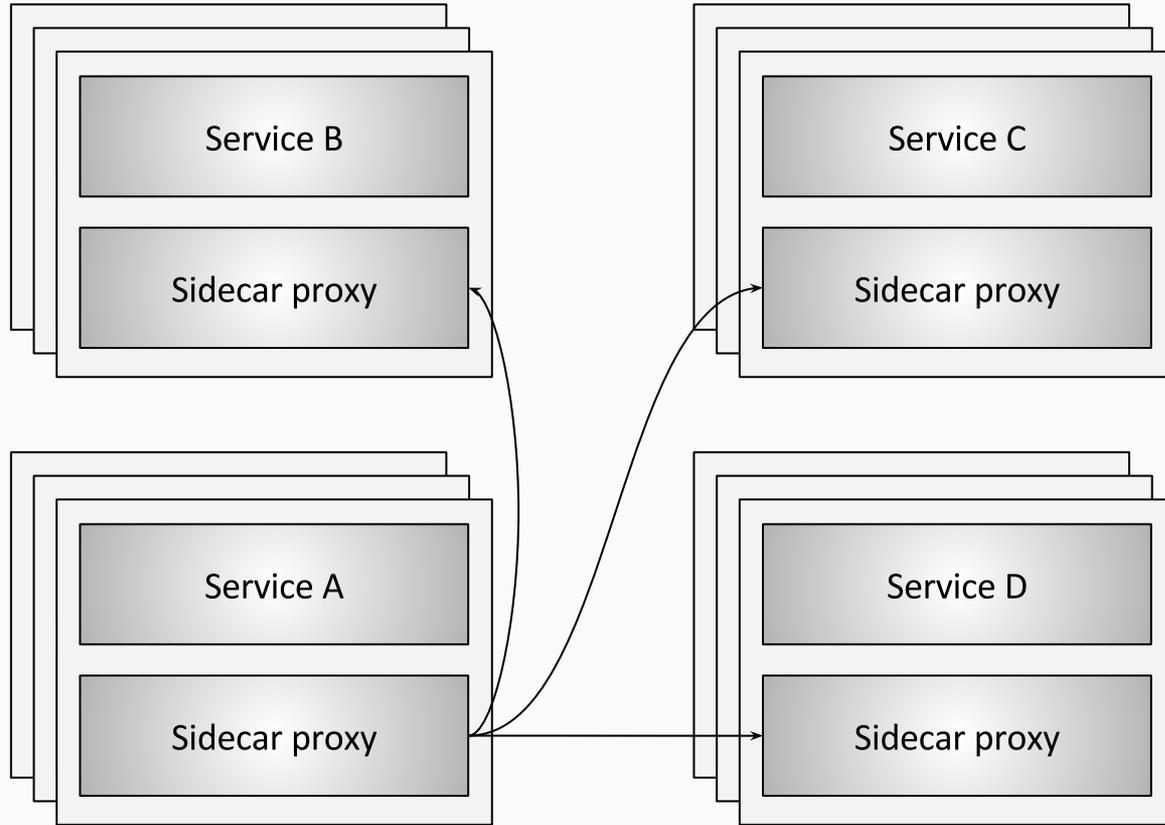
## What is Envoy and the service mesh?

*The network should be transparent to applications. When network and application problems do occur it should be easy to determine the source of the problem.*



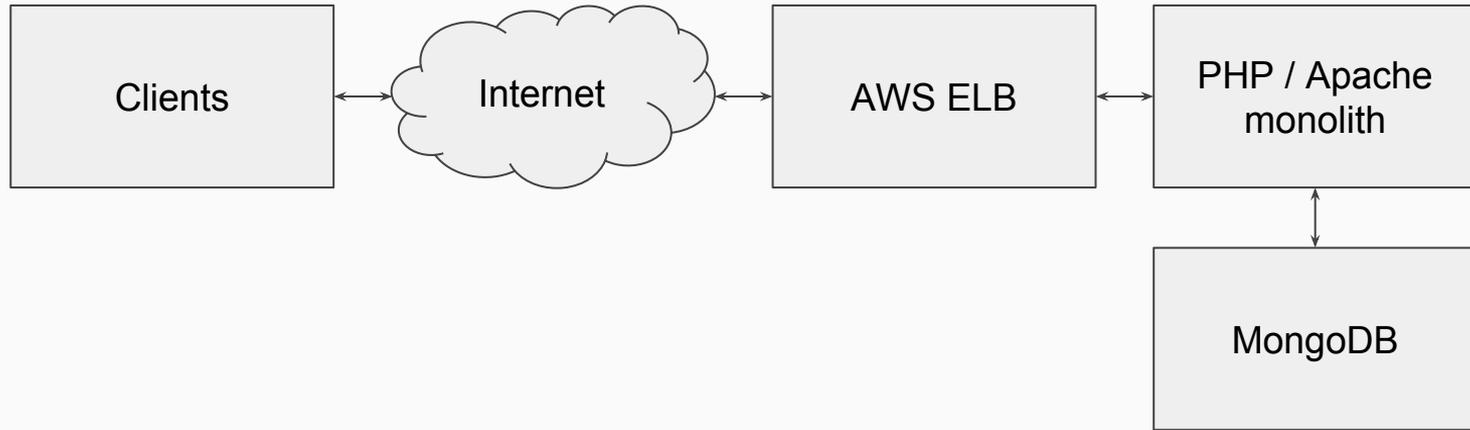
envoy

# Service mesh refresher



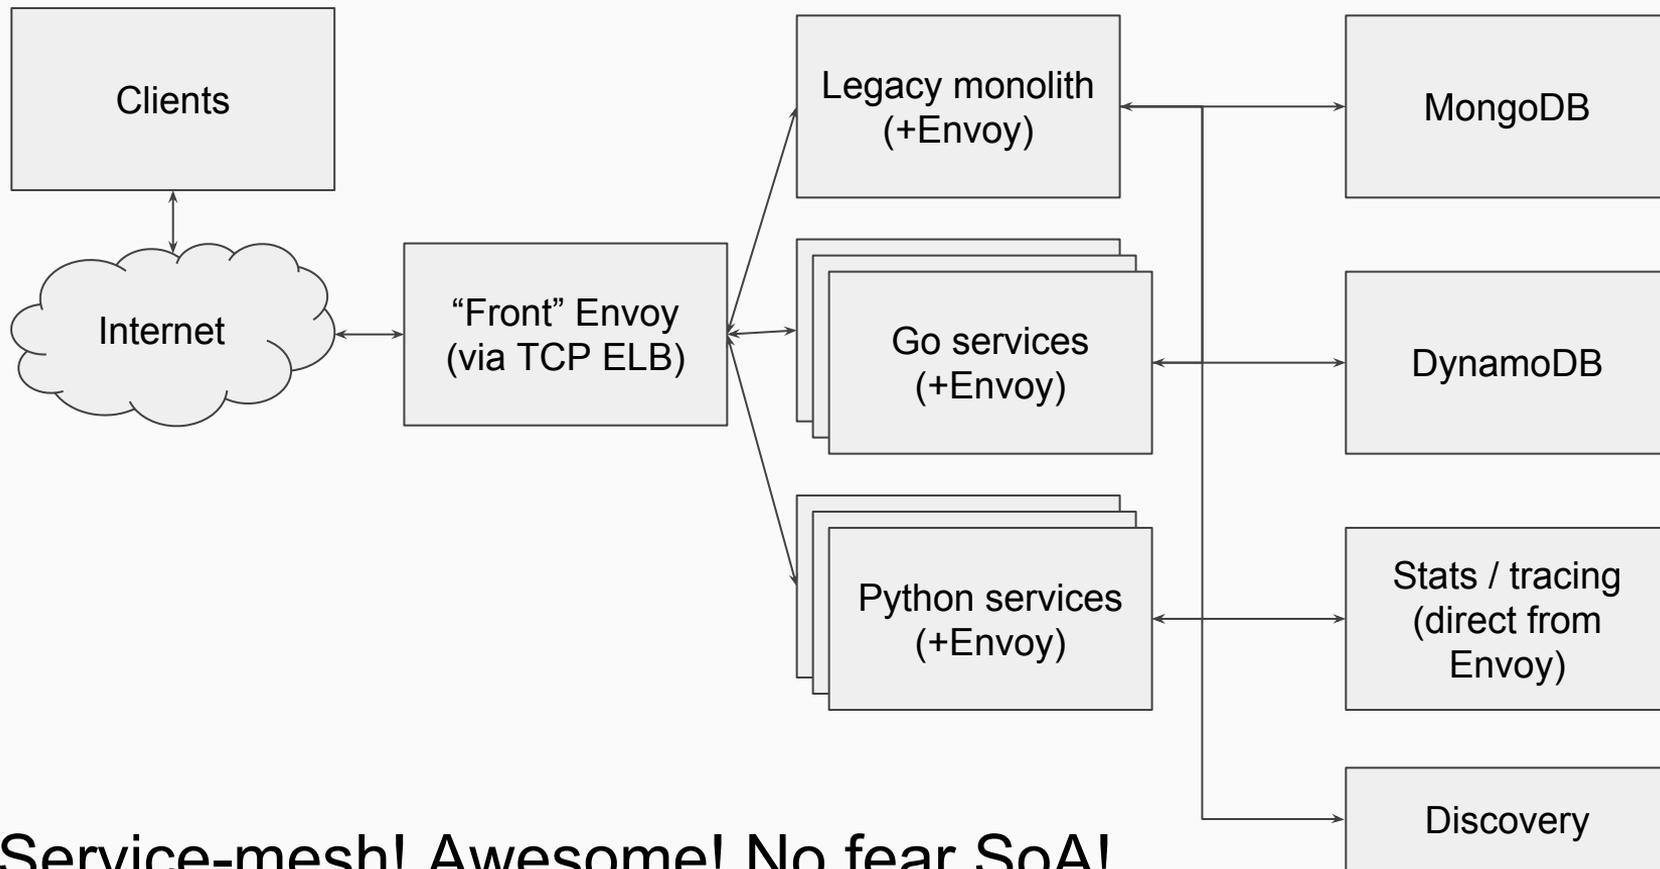
- **Out of process architecture**
- **Modern C++11 code base**
- **L3/L4 filter architecture**
- **HTTP L7 filter architecture**
- **HTTP/2 first**
- **Service discovery and active/passive health checking**
- **Advanced load balancing**
- **Best in class observability** (stats, logging, and tracing)
- **Edge proxy**

Lyft ~4 years ago



Simple! No SoA! (*but still not that simple*)

# Lyft today



**Service-mesh! Awesome! No fear SoA!**

## How did Lyft go from then to now

- Can't go from before to after overnight.
- Must be incremental. Show value in steps.
- Perfect is the enemy of the good.
- This is how we did it...(teaser: it wasn't easy)





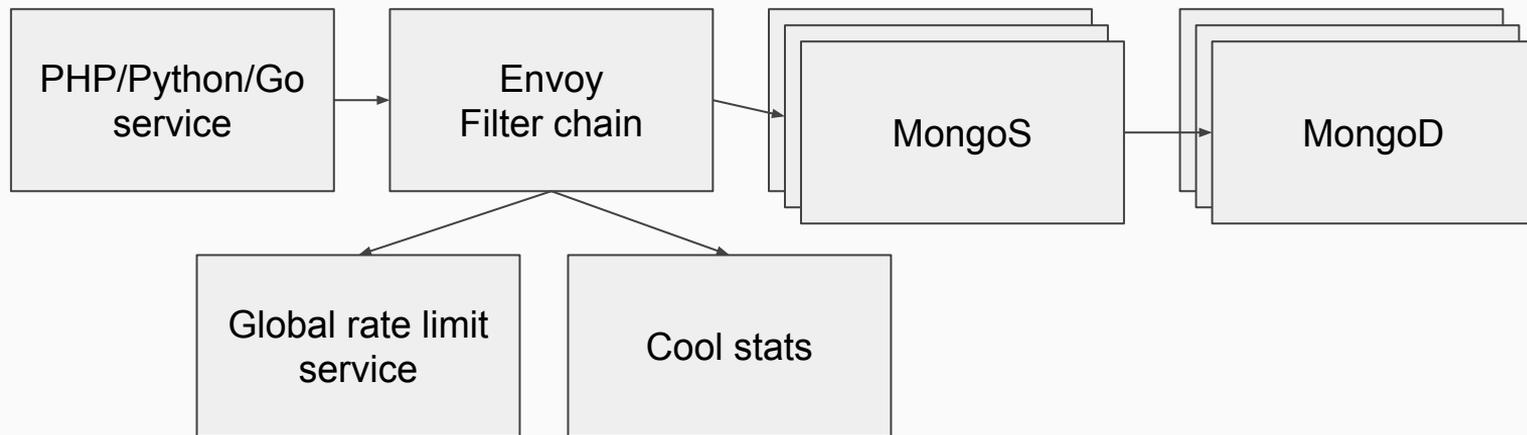
# Start with edge proxy

Observability, observability, observability...



## Next: TCP proxy and Mongo

- PHP to sharded Mongo not efficient with connection counts.
- Mongo bad at connection handling.
- Limit connections into Mongo.
- ... We can parse Mongo at L7 and generate cool stats!
- ... We can ratelimit Mongo to avoid death spirals!
- ... We can do this for all services, not just Python!



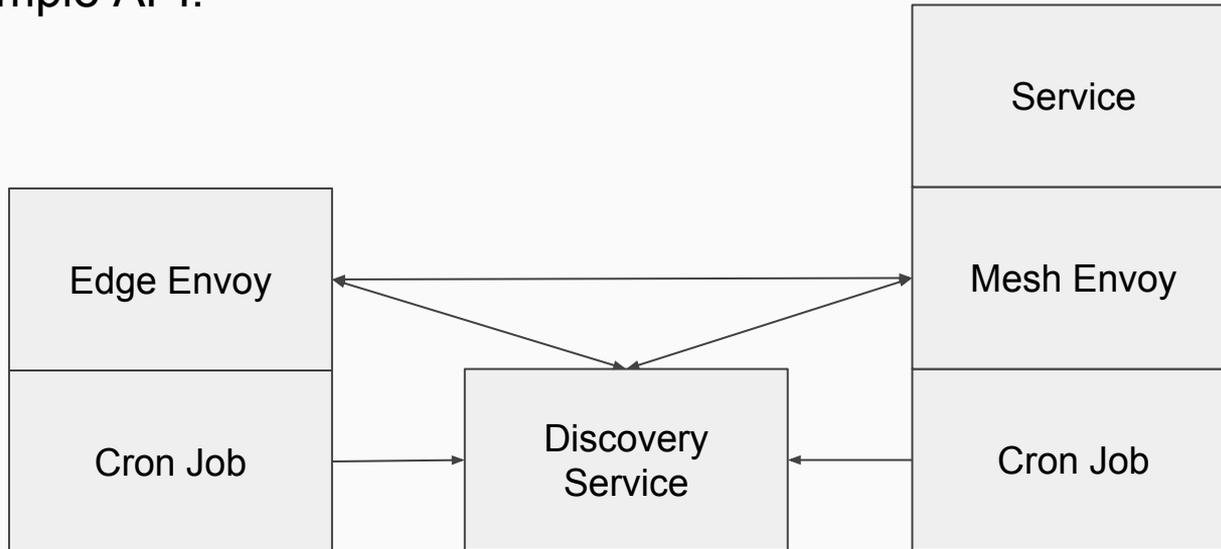
## Next: service sidecar

- Because of Mongo, Envoy already running on services.
- Let's use for ingress buffering, circuit breaking, and observability.
- Still using internal ELBs at this point for service to service traffic.
- Still get tons of value out of above without direct mesh connect.



## Next: service discovery and real mesh

- ELBs are (and intermediate LBs in general) ... not great. Let's do direct connect.
- Need service discovery.
- No ZK, etcd, or Consul. Eventually consistent. Build dead simple system with dead simple API.



# Envoy thin clients @Lyft

```
from lyft.api_client import EnvoyClient
switchboard_client = EnvoyClient(
    service='switchboard'
)
msg = {'template': 'breaksignout'}
headers = {'x-lyft-user-id': 12345647363394}
switchboard_client.post("/v2/messages", data=msg, headers=headers)
```

- Abstract away egress port
- Request ID/tracing propagation
- Guide devs into good timeout, retry, etc. policies
- Similar thin clients for Go and PHP

## Next ... Envoy everywhere

- OK this Envoy thing is pretty neat.
- Need to run it *everywhere* for full value.
- And so begins the slog. Many month burndown to convert everything and remove ELBs.
- Why a slog? Because of how Lyft manages services and configurations...
- But after slog complete, payoff is tremendous. Keep adding *features that developers want* and deploy them widely very quickly...



## Initial Envoy config management

- Envoy config is JSON/YAML.
- Initially we had full JSON committed.
- Deploys compile Envoy and bundle configs. Don't need to do back compat.
- Config per deployment type.
  - front\_envoy.json
  - service\_to\_service\_envoy.json
- Envoy deployed via “pull deploy” and salt on each host.
- 1 service per host, 1 envoy per service, 1 envoy per host.
- Hot restart used for both config/binary deploys (no difference).

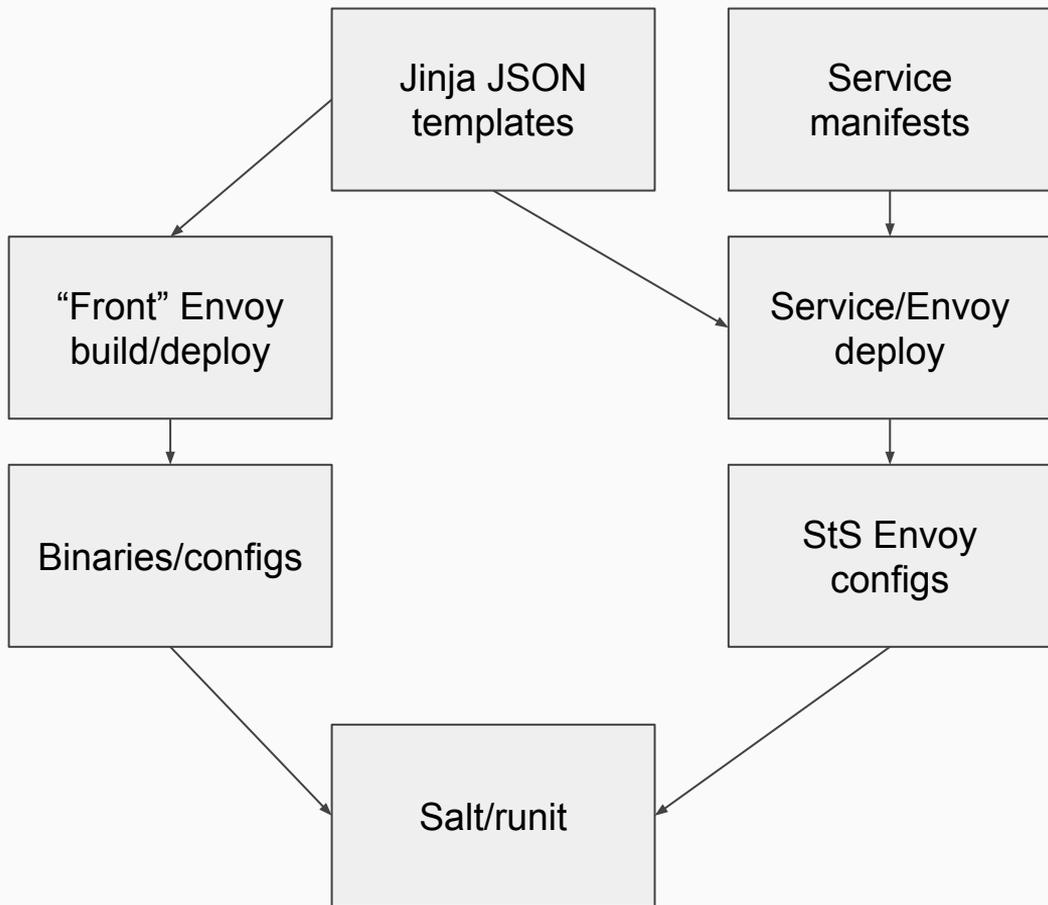
```
{  
  "listeners": {...},  
  "clusters": {...},  
  ...  
}
```

## Next comes configgen.py and ... jinja!

- These configs are starting to get really tedious to write!
- Would like some amount of static scripting.
- Develop a tool called configgen.py which takes in a bunch of templates and inputs, runs jinja on them, and produces final outputs.
- Already starting to see trend that ultimately no way around complex Envoy configs being **machine generated**.

```
{% import 'certs.json' as certs -%}
{% macro listener(port,ssl,proxy_proto) %}
{
  "address": "tcp://0.0.0.0:{{ port }}",
  {% if ssl -%}
  "ssl_context": {
    "alpn_protocols": "h2,http/1.1",
    {{ certs.public(service_instance) }}
  },
  {% endif -%}
}
```

# Envoy config/process management @Lyft with templating

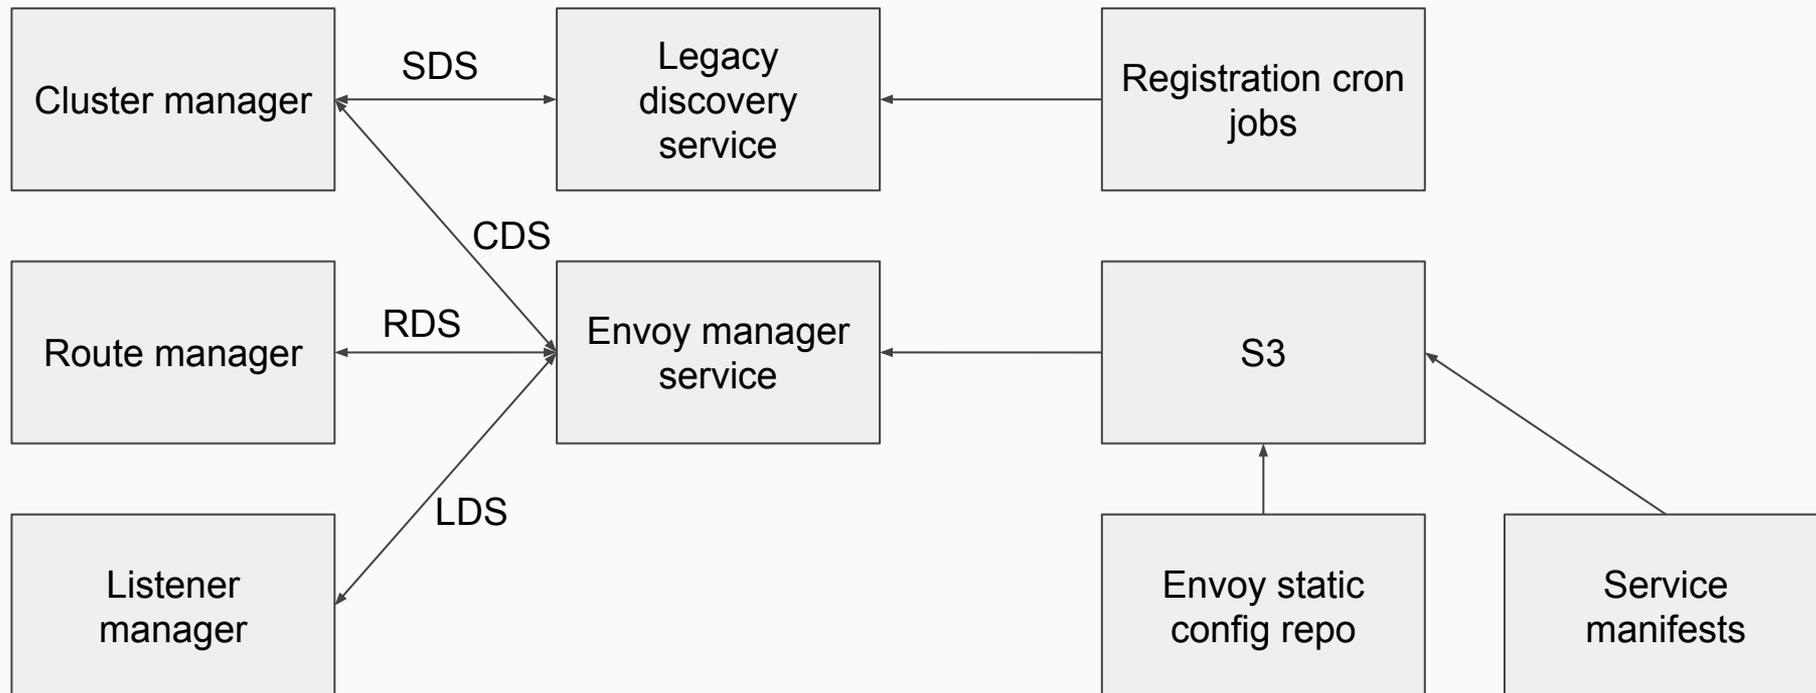


- Combination of static and dynamic configs.
- Service egress, circuit breaking, etc. configs specified in manifest.
- Service configs built on service host at service/envoy deploy time.

## Envoy control plane APIs

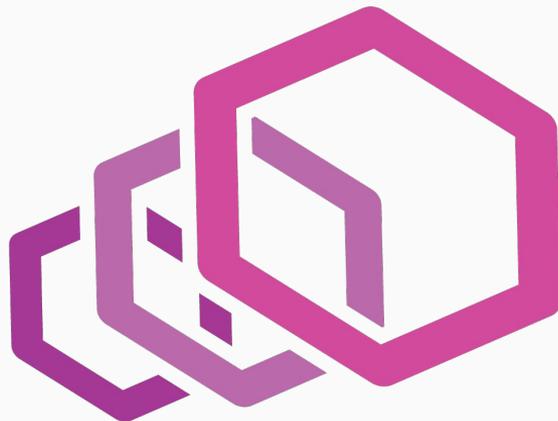
- The goal of Envoy has always been as a *universal data plane*.
- Support APIs to enable control plane integration.
- V1 JSON/REST APIs (order of implementation):
  - **Service** Discovery Service (SDS) (note: poorly named)
  - **Cluster** Discovery Service (CDS)
  - **Route** Discovery Service (RDS)
  - **Listener** Discovery Service (LDS)
- @Lyft we **only currently use SDS**. Everything else driven by Jinja/JSON templating.

# Lyft Envoy config management now



Only need a very tiny bootstrap config for each envoy...

- Thanks for coming! Questions welcome on Twitter: **@mattklein123**
- We are super excited about building a **community** around Envoy. Talk to us if you need help getting started.
- <https://www.envoyproxy.io/>
- **Lyft is hiring!**

The Lyft logo is displayed in white text on a solid magenta square background. The word "Lyft" is written in a rounded, lowercase sans-serif font.

envoy